

Online Prediction of Dyadic Data with Heterogeneous Matrix Factorization

Guangyong Chen * Fengyuan Zhu * Pheng Ann Heng *

January 14, 2016

Abstract

Dyadic Data Prediction (DDP) is an important problem in many research areas. This paper develops a novel fully Bayesian nonparametric framework which integrates two popular and complementary approaches, discrete mixed membership modeling and continuous latent factor modeling into a unified Heterogeneous Matrix Factorization (HeMF) model, which can predict the unobserved dyadics accurately. The HeMF can determine the number of communities automatically and exploit the latent linear structure for each bicluster efficiently. We propose a Variational Bayesian method to estimate the parameters and missing data. We further develop a novel online learning approach for Variational inference and use it for the online learning of HeMF, which can efficiently cope with the important large-scale DDP problem. We evaluate the performance of our method on the EachMoive, MovieLens and Netflix Prize collaborative filtering datasets. The experiment shows that, our model outperforms state-of-the-art methods on all benchmarks. Compared with Stochastic Gradient Method (SGD), our online learning approach achieves significant improvement on the estimation accuracy and robustness.

1 Introduction

Over the past few decades, dyadic data prediction (DDP) problem [1] has attracted lots of research attentions from many areas, including pattern recognition, neural computation, statistics, machine learning and data mining. Dyadic means an ordered pair of objects and the goal of DDP is to predict the value of unseen dyadic given the values of observed ones. Mathematically, Candes and Tao [2] defined the DDP problem as the process of “recovering the whole matrix $R \in \mathbb{R}^{U \times M}$ from only a sampled set of its entries $\{r_{ij}\}_{(i,j) \in \Omega}$, with Ω being a subset of complete set of entries $[U] \times [M]$ ”. Here and in the sequel, $[U]$ denotes the set $\{1, \dots, U\}$. As illustrated in Fig. 1, the task of DDP is to estimate the values of white entries (unobserved dyadic) in each matrix given the colored entries denoting the observed dyadic. An important application of this method is the recommendation system which predicts the preference of users to a special item based on the learned users’ taste information and item’s latent properties. This problem also finds large amount of practical applications in other research fields, including the image completion task in the filed of computer vision, blind affinity prediction in Bioinformatics [3], click prediction in Web search [4], and other applications [5, 6, 7, 8, 9].

The DDP is an ill-posed problem and it is impossible to determine the missing entries without making any assumption about the matrix R . A popular approach for this problem is to assume that the unknown matrix R has low rank L or has approximately low rank L . It has been proved that under certain incoherence assumptions on the singular vectors of the matrix R , exact prediction is possible by solving a convenient convex program as long as the number of samples is on the order of $NL \log(N)$ with $N = \max\{U, M\}$ [10, 2, 11]. These seminal results have boomed the development of this filed. However, in some practical situations, DDP still remains intractable when it is impossible to acquire enough entries to satisfy the minimum number condition. This problem is common in the area of recommendation system. For example, in the dataset of *MovieLens 100k*, the rating matrix R is highly sparse, where only around 5% entries are

*All authors are with the Chinese University of Hong Kong, Shatin, N.T., Hong Kong. The contact information can be found on the website <http://www.cse.cuhk.edu.hk/~pheng/>

observed. To tackle this issue, [12, 13, 14] proposed the Bayesian Probabilistic Matrix Factorization (BPMF) methods which model R as the product of an user coefficient matrix $A \in \mathbb{R}^{L \times U}$ and an item coefficient matrix $B \in \mathbb{R}^{L \times M}$, with both matrices A and B following Gaussian distributions.

Though BPMF methods perform remarkably well on the DDP task, they fail to capture the heterogeneous nature of objects and their interactions because of the Gaussian assumption. The heterogeneous natural of objects is common in real commercial recommendation systems. For example, in *Douban.com*, users are usually clustered into different communities depending on their own interests, and items are often categorized into multiple groups based on their own properties. Thus, it is not surprising that users from different groups may have quite distinct opinions for some movies, such as "*Napoleon Dynamite*", "*Kill Bill: Volume 1*", "*Sideways*". The Gaussian assumption of BPMF cannot capture this feature properly resulting to poor prediction results. This issue is significant and should be well considered when developing algorithms for solving DDP tasks, or the prediction performance will be highly affected empirically. To tackle this issue, models like Mixed Membership Stochastic Blockmodel [15] and Bi-LDA [16] have been proposed with the assumption that objects are generated from different communities. They introduce context dependence by allowing each object to select a new topic for each new interaction. However, the relatively poor predictive performance of Bi-LDA suggests that the blockmodel assumption is still too restrictive. This paper proposes a novel Heterogeneous Matrix Factorization (HeMF) model to unify discrete mixed membership model with BPMF, where the missing entries are estimated from the divided homogeneous sub-matrices. Fig. 1 (b) illustrates the HeMF for better understanding. The proposed HeMF model can be also viewed as a new type of bicluster model, where each bicluster have an intrinsic linear structure. Compared with the traditional Matrix Factorization approaches [17, 18, 19, 20, 21], the proposed model HeMF incorporates the community membership information from a new perspective. Because the number of bi-clusters are usually unknown, this work further incorporates the Bayesian nonparametric technique and let the data determine the model complexity automatically. This paper further derives an efficient batch-inference algorithm for HeMF under the principle of Variational Bayesian (VB) [22]. As demonstrated empirically, the derived method converges much faster than traditional sampling methods, and gives better prediction performance.




Figure 1: An illustration for the rating matrix R evolved with time, where I_1 means the 1st item and U_1 represents the 1st user. For objects from different communities, their index are marked by different colors accordingly. (a) denotes the sparse rating matrix acquired at time t . (b) contains 4 homogeneous sub-matrices divided by the rating matrix observed at time t . (c) denotes the rating matrix observed at time $t+1$, where new observations are marked by the blue color. Note that the 2nd user changed his rating score for the 7st item. (d) contains the rating matrix observed at time $t+2$, where new observations are marked by the color green, and new item and user are registered.

The era of *Big Data* presents new challenges for our DDP task. Many real world applications involve massive amount of data that even cannot be accommodated entirely in the memory. Moreover, the amount of data even increases without any limitation as time goes on. For example, users give new ratings to items at nearly all time in Amazon, and there are always new user accounts registered online and new items launched into the market. Mathematically speaking, new entries are labeled in the rating matrix R

continuously, and the size of R will become larger as time goes on. This process can be better understood in Fig. 1. Because both sampling and variational methods maintain the entire configuration and perform iterative updates of multiple passes, these inference strategies limit their implementations on the massive stream data. Stochastic Gradient Descent (SGD) and sampling approaches have been applied in a sequential setting for matrix factorization [23, 24]. However, they are not embarrassingly online and hence cannot be directly applied to the stream data. As observed empirically, SGD always gives the results vibrating around an unsatisfied solution and faces a serious overfitting problem. These challenges motivate us to develop a novel principle of variational Bayesian for massive stream dataset, which is different from the pioneer works [25, 26, 27, 28]. Moreover, as our best knowledge, our approach is the first online variational inference approach to the problem of DDP task.

We summarize the contributions of this work as follows:

- We develop a new Heterogeneous Matrix Factorization (HeMF) model for DDP task, which can efficiently capture the heterogeneous natural for Dyadic data in recommender systems. We construct HeMF in a Bayesian nonparametric manner and learn the number of communities and hidden dimensionality automatically during the implementation. HeMF model can be also viewed as a novel type of Bi-Cluster model, and it is the first method to introduce intrinsic linear structures into each bi-cluster.
- We develop a batch VB (bVB) inference method for learning the proposed HeMF model, which achieves better performance and converges faster compared with the state-of-the-art methods.
- We investigate an emerging problem in our paper, which is *how to handle the massive data problem in DDP task*. This is an interesting problem in the field of DDP especially in the application of recommender system. Under the principle of Variational Bayesian, we derive a novel online Variational Bayesian (oVB) method to tackle this problem in a fraction of the time required by traditional inference. The performance of oVB is naturally guaranteed.
- We further develop an empirical Variational Bayesian (eVB) procedure to improve the performance of bVB and oVB. We apply both algorithms on real Dyadic data and the experimental results demonstrate that our methods significantly outperform the state-of-the-art approaches consistently.

We organize the rest of this paper as follows, Sec. 2 reviews the background of Dirichlet Process for the construction of HeMF model; Sec. 3 introduces our HeMF model which is a Coupled Dirichlet Process Model to describe the generative process of factor vectors; Sec. 4 develops an efficient bVB inference method to infer the proposed HeMF model with batch of data; Sec. 5 derives a novel online learning method under a newly proposed online learning principle; Sec. 6 further proposes an empirical Variational Bayesian (eVB) to infer the hyper-parameters for bVB and oVB algorithm to improve their performance; to evaluate the performance of our method, Sec. 7 proposes extensive comparative studies of our approach with previous ones on four real-world dataset, where our methods achieve superior performance over the competitive ones; and Sec. 8 concludes this paper.

2 Dirichlet Process Mixture Model

This section introduces the Dirichlet Process Mixture Model (DPMM), which is among the most popular clustering models in practice for analyzing the heterogeneous data. Different from the traditional parametric models for clustering, DPMM allows the number of groups to vary during inference, which provides great flexibility for exploratory analysis.

We first introduce the Dirichlet Process (DP) [29], which is typically denoted by $DP(\alpha, \mu)$ with a concentration parameter α and a base distribution μ . A DP can be well constructed via the Chinese Restaurant Process (CRP). Given a Chinese restaurant with countably infinite tables, customers walk in one after another and sit down at a certain table with the following scheme:

1. The first customer always chooses the first table.
2. The t th customer chooses either an unoccupied table with probability $\frac{\alpha}{t + \alpha}$; or an occupied table with probability $\frac{c}{t + \alpha}$, where c is the number of people sitting at that table.

For the convenience of inference, Sethuraman [30] proposed a stick-breaking approach for DP construction, which can be defined as

$$\begin{aligned}
D &= \sum_{k=1}^{\infty} \pi_k \delta_{\phi_k}, \\
\text{with } \phi_k &\sim \mu, \forall k = 1, 2, \dots, \\
\pi_k &= v_k \prod_{l=1}^{k-1} v_l, v_k \sim \text{Beta}(1, \alpha).
\end{aligned} \tag{1}$$

Here, $\text{Beta}(\alpha, \beta)$ denotes a Beta distribution with parameter α and β .

The sample paths of a DP are almost sure discrete. Due to this nice property, the DP is widely used in the construction of mixture model which is very useful for modeling heterogeneous data. A DPMM can be well expressed with the following generative process

$$\begin{aligned}
D &\sim DP(\alpha, \mu), \\
\theta_i &\sim D, \\
x_i &\sim F(\cdot | \theta_i), \forall i = 1, 2, \dots, n,
\end{aligned} \tag{2}$$

where data x_1, \dots, x_n are the realizations from distribution F with parameter $\theta_1, \dots, \theta_n$. Because an atom ϕ_k can be repeatedly generated from D with positive probability, there is a partition $\{\mathcal{P}_1, \dots, \mathcal{P}_K\}$ of $\{1, \dots, n\}$ such that θ_i are identical for all $i \in \mathcal{P}_k$, which we denote by ϕ_k .

The DPMM serves as a foundation for varieties of Bayesian nonparametric models, and has achieved substantial progress on representing feature-based data. We will use the DP and DPMM to construct our HeMF model for the DDP task.

3 Heterogeneous Matrix Factorization with Coupled Dirichlet Process

This section describes the Heterogeneous Matrix Factorization (HeMF) model for DDP, which utilizes a coupled Dirichlet Process to describe the generative model of a rating matrix. Suppose we have U users and M items, and let r_{ij} be the rating of user i for item j , then we have:

$$r_{ij} = a_i^T b_j + e_{ij}. \tag{3}$$

Here, $b_j \in \mathbb{R}^{L \times 1}$ denotes an item-specific feature vector, where each element represents an objective score for one criterion, such as picture, directing, actor, actress, etc used in Academy Award of Merit; and $a_i \in \mathbb{R}^{L \times 1}$ denotes the preference of user i for these criterion. Thus, $a_i^T b_j$ presents the weighted average score of user i for item j . Moreover, e_{ij} denotes the zero-mean Gaussian noise with variance σ^2 , which captures the uncertainty of rating behavior. Thus, the conditional distribution of the rating matrix $R \in \mathbb{R}^{U \times M}$ over A and B is given by

$$p(R_\Omega | A, B, \sigma^2) = \prod_{(i,j) \in \Omega} \mathcal{G}(r_{ij} | a_i^T b_j, \sigma^2), \tag{4}$$

where $\mathcal{G}(x | \mu, \Sigma)$ denotes the Gaussian distribution with the mean vector μ and the covariance matrix Σ , and Ω denote the given observed index set .

Traditional BPMF approaches are commonly developed with the assumption that a_i and b_j are Gaussian distributed. However, this assumption is not suitable for real commercial recommendation systems because of the heterogeneous natural of users and items. For example, in the commercial recommendation system *Moivlens*, a movie can be categorized as *computer animation, dramatic, touching or other labels*, while a user can be labeled by their specialties, such as *historian, scientist, engineer, poet*. It is obvious that users from different specialties have their own preference for movies. Thus, the BPMF with Gaussian assumption is unsuitable for DDP task in the problem. It is significant to model such heterogeneous property for better

prediction result. To tackle this issue, the HeMF model introduces the DPMM to model the a_i and b_j respectively resulting to a novel coupled Dirichlet Process to capture this heterogeneous natural.

The motivation of applying Dirichlet Process in our paper is very intuitive and can be well explained from the perspective of CRP which is one of the construction of DP. Imaging a recommendation system with infinite number of interest communities, each with infinite capacity. The 1^{st} user construct his own interest community with probability 1. At time $t + 1$, a new user chooses at random to participate in one of the following $D + 1$ interest communities: directly to the d -th already constructed community with probability $\frac{|d|}{t+1+\alpha}$ where $|d|$ is the size of d -th community, or establishes a new one with the probability $\frac{\alpha}{t+1+\alpha}$.

For the convenience of inference, this paper uses the stick-breaking construction to build the HeMF and its generative process is as follows:

- For i -th user
 - Draw $\pi_d \sim \text{Beta}(1, \alpha)$
 - Draw $\phi_d \sim H$
 - Draw $z_i = d \sim \pi_d \prod_{t=1}^{d-1} (1 - \pi_t)$
 - Draw $a_i \sim F(\cdot | \phi_{z_i})$
- For j -th item
 - Draw $\omega_k \sim \text{Beta}(1, \beta)$
 - Draw $\psi_k \sim G$
 - Draw $\tilde{z}_j = k \sim \omega_k \prod_{t=1}^{k-1} (1 - \omega_t)$
 - Draw $b_j \sim F(\cdot | \psi_{\tilde{z}_j})$
- Sample the rating value of their interaction
 - Draw $r_{ij} \sim \mathcal{G}(\cdot | a_i^T b_j, \sigma^2)$.

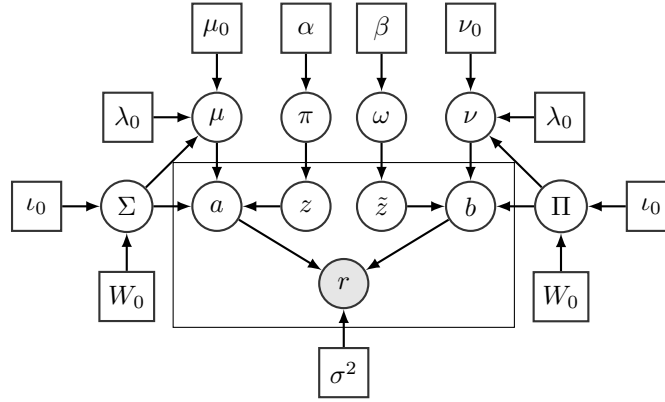


Figure 2: The graphical model for illustrating the proposed HeMF with Coupled Dirichlet Process, where the shadow node means the observed rating score, and the nodes covered by rectangle box mean the missing data. The leaf nodes illustrated in the graphical model mean the hyperparameters of the proposed HeMF model.

It can be observed that there are two separate Dirichlet Processes for each user-item interaction, which can be better explained in Fig. 2. Given D communities for users and K groups for items in total, an observed rating matrix R is divided into $D \times K$ homogeneous sub-matrices, as shown in Fig. 1. The HeMF efficiently combine the division step and the matrix factorization step into a unified framework.

3.1 Constructing Base Distributions

Given the HeMF model, we further propose the base distribution H and G which should capture the previously mentioned low-rank properties. With our approach, the latent feature vectors from the same community should surround a low-rank space. Before stepping into the detailed constructions of H and G , we first introduce the following lemma to construct a proxy matrix Y for the target matrix X , whose columns coming from the same community.

Lemma 1. *For a matrix $X \in \mathbb{R}^{m \times n}$ with $m \leq n$, its rank is bounded by the transferred matrix $Y = W^{\frac{1}{2}}(X - M)$, where $W \in \mathbb{R}^{m \times m}$ is a full-rank semi-definite positive matrix and each column in the matrix $M \in \mathbb{R}^{m \times n}$ is the mean of all columns in the matrix X .*

Proof. Because each column in the matrix M is the same with each other, we have $\text{rank}(M) = 1$. And we have $\text{rank}(X - M) \leq \text{rank}(X)$ because all columns in both X and M can be represented by the columns in X . With the sub-additivity property of rank operation, we have

$$\text{rank}(X) \leq \text{rank}(X - M) + \text{rank}(M). \quad (5)$$

Thus, we can conclude that

$$\text{rank}(X) - 1 \leq \text{rank}(X - M) \leq \text{rank}(X). \quad (6)$$

Because W is a full-rank matrix, we achieve

$$\text{rank}(X) - 1 \leq \text{rank}(Y) \leq \text{rank}(X), \quad (7)$$

which means that the rank of the target matrix is bounded by the following inequalities,

$$\text{rank}(Y) \leq \text{rank}(X) \leq \text{rank}(Y) + 1. \quad (8)$$

□

Thus, we can regularize the rank of the proxy matrix Y , instead of constraining the targeted matrix X directly. An advantage of introducing the proxy matrix into our model stems from the facts that Y provides much more flexibility when modeling X .

The direct optimization with rank regularization is clearly a NP-hard problem. Fortunately, as proved in [31], the minimum rank approximation can be achieved by minimizing the nuclear norm. Thus, we can obtain,

$$\min \text{rank}(X) \propto \min \text{rank}(Y) \propto \min \text{tr}(Y^T Y). \quad (9)$$

This approach has been widely used in machine learning community to promote a low rank solution without a pre-define L , and can be further formulated as

$$\begin{aligned} \min \text{tr}(Y^T Y) &\propto \min \sum_{t=1}^n (x_t - m)^T W (x_t - m) \\ &\propto \max -\frac{1}{2} \sum_{t=1}^n (x_t - m)^T W (x_t - m) \\ &\propto \max \ln \prod_{t=1}^n \mathcal{G}(x_t | m, W^{-1}) \end{aligned} \quad (10)$$

with $m = \frac{1}{n} \sum_{t=1}^n x_t$. Thus, it can be found that regularizing the rank of X is equivalent to assuming each column of the matrix to follow a Gaussian distribution. Thus, following [14], for each user-item interaction, we assume a_i and b_j to be drawn from the following Gaussian distributions,

$$\begin{aligned} F(\cdot | \phi_d) &= \mathcal{G}(\cdot | \mu_d, \Sigma_d), \\ F(\cdot | \psi_k) &= \mathcal{G}(\cdot | \nu_k, \Pi_k), \end{aligned} \quad (11)$$

with $\phi_d = \{\mu_d, \Sigma_d\}$ and $\psi_k = \{\nu_k, \Pi_k\}$. As in [14, 17], we consider a fully Bayesian treatment where the topic model parameters ϕ_d and ψ_k are the random realizations from the following distributions respectively,

$$\begin{aligned} H(\phi_d) &= \mathcal{G}(\mu_d | \mu_0, \lambda_0 \Sigma_d) i\mathcal{W}(\Sigma_d | W_0, \iota_0), \\ G(\psi_k) &= \mathcal{G}(\nu_k | \nu_0, \lambda_0 \Pi_k) i\mathcal{W}(\Pi_k | W_0, \iota_0), \end{aligned} \quad (12)$$

where $i\mathcal{W}(\cdot)$ denotes the inverse Wishart distribution.

Thus, we can obtain the marginal distribution of sparsely observed rating matrix R as follows:

$$p(R_\Omega | \Xi) = \int p(R_\Omega, Y, \Theta | \Xi) dY d\Theta, \quad (13)$$

where $\Xi = \{\mu_0, \nu_0, \sigma^2, W_0, \iota_0, \lambda_0, \alpha, \beta\}$ denotes the hyper-parameters of the proposed model, $\Theta = \{\phi_d\}_{d=1}^D \cup \{\psi_k\}_{k=1}^K \cup \{\pi, \omega\}$ denotes the model parameters, and $Y = \{a_i, z_i\}_{i=1}^U \cup \{b_j, \tilde{z}_j\}_{j=1}^M$ denotes the missing data associated with each observed entry in the rating matrix R . The relationships between parameters has been illustrated in Fig. 2, where the observed rating r is shown by a shaded node and the missing variables associated with the observed rating r is covered by a rectangle box. The leaf nodes of this graphical model is shown by a rectangle node and denotes the a hyper-parameter of HeMF. Thus, $p(R_\Omega, Y, \Theta | \Xi)$ can be decomposed as follows,

$$\begin{aligned} p(R_\Omega, Y, \Theta | \Xi) &= p(R_\Omega | A, B, \sigma^2) p(A | Z, \mu, \Sigma) p(B | \tilde{Z}, \nu, \Pi) \\ &\quad \cdot p(Z | \pi) p(\pi | \alpha) \prod_{d=1}^D \mathcal{G}(\mu_d | \mu_0, \lambda_0 \Sigma_d) i\mathcal{W}(\Sigma_d | W_0, \iota_0) \\ &\quad \cdot p(\tilde{Z} | \omega) p(\omega | \beta) \prod_{k=1}^K \mathcal{G}(\nu_k | \nu_0, \lambda_0 \Pi_k) i\mathcal{W}(\Pi_k | W_0, \iota_0), \end{aligned}$$

with

$$\begin{aligned} p(Z | \pi) &= \prod_{i=1}^U \pi_{z_i} \prod_{t=1}^{z_i-1} (1 - \pi_t), \quad p(\pi | \alpha) = \prod_{d=1}^D \text{Beta}(\pi_d | 1, \alpha), \\ p(\tilde{Z} | \omega) &= \prod_{j=1}^M \omega_{\tilde{z}_j} \prod_{t=1}^{\tilde{z}_j-1} (1 - \omega_t), \quad p(\omega | \beta) = \prod_{k=1}^K \text{Beta}(\omega_d | 1, \beta), \\ p(A | Z, \mu, \Sigma) &= \prod_{i=1}^U \mathcal{G}(a_i | \mu_{z_i}, \Sigma_{z_i}), \\ p(B | \tilde{Z}, \nu, \Pi) &= \prod_{j=1}^M \mathcal{G}(b_j | \nu_{\tilde{z}_j}, \Pi_{\tilde{z}_j}). \end{aligned}$$

As shown in Eq. 13, the Bayesian framework considers a whole class of models, rather than focusing on a single model to provide a solution to the DDP task. Thus, the HeMF with Bayesian framework can avoid overfitting problem by integrating out the parameters.

4 Variational Inference

Unfortunately, direct marginal inference of Eq. 13 is intractable. Markov chain Monte Carlo methods [17] have been widely used to achieve exact marginal results, but typically require vast computational resources and become inefficient for complex models in high data dimensions. In this section, we derive a batch-based variational Bayesian inference method, which is a practical framework for Bayesian computations in graphical models.

Given the rating matrix R observed on a sampled set of entries Ω , Variational Bayesian (VB) approach introduces a trial distribution $q(Y, \Theta)$ to maximize the lower bound of the marginal distribution, $\mathcal{L}_{q(Y, \Theta)}(R_\Omega | \Xi)$, as follows,

$$\ln p(R_\Omega | \Xi) \geq \mathcal{L}_{q(Y, \Theta)}(R_\Omega | \Xi), \quad (14)$$

where

$$\mathcal{L}_{q(Y, \Theta)}(R_\Omega | \Xi) = \text{KL}(q(Y, \Theta) || p(R_\Omega, Y, \Theta | \Xi)). \quad (15)$$

Here, $\text{KL}(\cdot)$ is the Kullback-Leibler (KL) distance between the trial distribution $q(Y, \Theta)$ and the joint distribution $p(R_\Omega | Y, \Theta, \Xi)p(Y, \Theta | \Xi)$. From Jensen's inequality, the lower bound of $\ln p(R_\Omega | \Xi)$ can be achieved by setting $q(Y, \Theta) \propto p(R_\Omega | Y, \Theta, \Xi)p(Y, \Theta | \Xi)$. In VB approximation, we often assume that $q(Y, \Theta)$ can be factorized as follows,

$$q(Y, \Theta) = q(Y)q(\pi)q(\omega) \prod_{d=1}^D q(\phi_d) \prod_{k=1}^K q(\psi_k), \quad (16)$$

where $q(Y)$, $q(\phi_d)$ and $q(\psi_k)$ can be further factorized as,

$$\begin{aligned} q(Y) &= \prod_{i=1}^U q(a_i)q(z_i | a_i) \prod_{j=1}^M q(b_i)q(\tilde{z}_i | b_i), \\ q(\phi_d) &= q(\mu_d)q(\Sigma_d | \mu_d), \quad q(\psi_k) = q(\nu_k)q(\Pi_k | \nu_k). \end{aligned}$$

This factorization corresponds to an approximation framework developed in physics called *mean field theory* [32], which breaks the entanglement among the model parameters, and leads to an efficient iterative algorithm.

Let θ_j denote the j -th parameter belonging to the set of parameters $Y \cup \Theta$. With other parameters fixed, the variational approximation problem with respect to $q(\theta_j)$ is equivalent to:

$$\ln p(R_\Omega | \Xi) \geq \int q(\theta_j) \ln \frac{\exp \mathbb{E}_{\neq j} \ln p(R_\Omega, Y, \Theta | \Xi)}{q(\theta_j)} d\theta_j, \quad (17)$$

where $\mathbb{E}_{\neq j}(\cdot)$ denotes an expectation with respect to the posterior distributions q over all parameters except θ_j . In this way, the close-form solution of $q(\theta_j)$ satisfies the following condition:

$$q(\theta_j) \propto \exp[\mathbb{E}_{\neq j} \ln p(R_\Omega | Y, \Theta, \Xi)p(Y, \Theta | \Xi)]. \quad (18)$$

Thus, the marginal distribution $q(R_\Omega | \Xi)$ can be solved by alternatively calculating Eq. 18 for each parameter.

Actually, the computations of the posterior distribution over θ_j can be greatly reduced by considering the graphical model, as Fig. 2. When updating variational approximation result with respect to θ_j , whose neighbor nodes are represented as $\theta_{\mathcal{N}_j}$, we can reformulate Eq. 18 as

$$q(\theta_j) \propto \exp[\mathbb{E}_{q(\theta_{\mathcal{N}_j})} \ln p(R_\Omega, \theta_j, \theta_{\mathcal{N}_j} | \Xi)], \quad (19)$$

where the number of involved parameters is dramatically decreased. For example, in the proposed HeMF model, the neighbor nodes of the stick-breaking vector π contain only the hyper-parameter α and the membership parameter z . Thus, we only need to consider the prior distribution $p(\pi | \alpha) \prod_{i=1}^U p(z_i | \pi)$, as well as the expectation with respect to the associated posterior distribution $\prod_{i=1}^U q(z_i)$ in the derivation of posterior distribution over π . Thus, we have

$$q(\alpha) \propto \exp[\mathbb{E}_{\prod_{i=1}^U q(z_i)} \ln p(\pi | \alpha) \prod_{i=1}^U p(z_i | \pi)]. \quad (20)$$

The detailed formulation of $q(\alpha)$ can be found in the later part of this section.

The updating formulas associated with the users' taste matrix A can be derived under the principle shown in Eq. 19, while the updating formulas associated with the items' property matrix B can be obtained by altering the notations. The parameters related to the users' taste matrix include A , Z , π and $\{\mu_d, \Sigma_d\}_{d=1}^D$. To simplify our formulations, let $\langle f(x) \rangle = \int q(x)f(x)dx = \mathbb{E}_{q(x)}f(x)$, which denotes the expectation of $f(x)$ with respect to the distribution $q(x)$.

Estimation of the users' taste matrix A : Based on the principle introduced in Eq. 19, we can obtain that the posterior distribution over a_i satisfies,

$$\begin{aligned} \ln q(a_i) \propto & -\frac{1}{2}a_i^T \left(\frac{\sum_{(i,j) \in \Omega} \langle b_j b_j^T \rangle}{\sigma^2} + \sum_{d=1}^D q_i(d) \langle \Sigma_d^{-1} \rangle \right) a_i \\ & + \left(\frac{\sum_{(i,j) \in \Omega} r_{i,j} \langle b_j^T \rangle}{\sigma^2} + \sum_{d=1}^D q_i(d) \langle \mu_d^T \rangle \langle \Sigma_d^{-1} \rangle \right) a_i. \end{aligned}$$

Since the parametric form of this posterior distribution is a quadratic function, the posterior distribution over a_i is Gaussian with the following parameters:

$$\begin{aligned} \langle a_i \rangle = & \left(\frac{1}{\sigma^2} \sum_{(i,j) \in \Omega} \langle b_j b_j^T \rangle + \sum_{d=1}^D q_i(d) \langle \Sigma_d^{-1} \rangle \right)^{-1} \\ & \cdot \left(\frac{1}{\sigma^2} \sum_{(i,j) \in \Omega} r_{i,j} \langle b_j \rangle + \sum_{d=1}^D q_i(d) \langle \Sigma_d^{-1} \rangle \langle \mu_d \rangle \right), \\ \langle a_i a_i^T \rangle = & \left(\frac{1}{\sigma^2} \sum_{(i,j) \in \Omega} \langle b_j b_j^T \rangle + \sum_{d=1}^D q_i(d) \langle \Sigma_d^{-1} \rangle \right)^{-1} \\ & + \langle a_i \rangle \langle a_i^T \rangle. \end{aligned} \tag{21}$$

From Eq. 21, we find that the expectation of the i -th user's taste is jointly determined by rated items and the information of belonged communities of the user. It can be understood from our practical experience that a human's taste can be not only reflected by the choices of items, but also the communities he or she belongs to.

Estimation of the membership vector of users' taste Z : After obtaining the updated users' taste matrix, we can get that the probability of each user's taste belonging to the d -th community satisfies

$$\ln q(z_i = d) \propto \gamma_i(d) = \xi_{1,i}(d) + \xi_{2,i}(d), \tag{22}$$

with

$$\begin{aligned} \xi_{1,i}(d) = & -\frac{1}{2}(\text{tr} \langle a_i a_i^T \rangle \langle \Sigma_d^{-1} \rangle - 2 \langle \mu_d^T \rangle \langle \Sigma_d^{-1} \rangle \langle a_i \rangle \\ & + \text{tr} \langle \mu_d \mu_d^T \rangle \langle \Sigma_d^{-1} \rangle + \langle \ln |\Sigma_d| \rangle), \\ \xi_{2,i}(d) = & \psi(\eta_{1,d}) - \psi(\eta_{1,d} + \eta_{2,d}) \\ & + \sum_{j=1}^{d-1} [\psi(\eta_{2,j}) - \psi(\eta_{1,j} + \eta_{2,j})], \end{aligned} \tag{23}$$

where $\psi(\cdot)$ denotes the digamma function [33], $\eta_{1,d}$ and $\eta_{2,d}$ denote the parameters to describe the posterior distribution over the d -th element in the vector π . Denoting $q(z_i = d)$ as $q_i(d)$, we can update the indicator variable z_i as follows:

$$q_i(d) = \frac{\exp(\gamma_i(d))}{\sum_{j=1}^D \exp(\gamma_i(j))}. \tag{24}$$

As observed in Eq. 22, the computation of $\gamma_i(d)$ consists of two terms, $\xi_{1,i}(d)$ and $\xi_{2,i}(d)$. Let $\xi_{2,i}(d) = 0$, we can find that $q_i(d)$ is the same with E-step in the classical Expectation-Maximization (EM) method [34].

Represented by the stick-breaking process, the effect of Dirichlet Process is introduced into our model by $\xi_{2,i}(d)$. The strategy of slightly revising the E-step to improve the performance of EM algorithm has been heuristically explored in some pioneer works, e.g. RPCL [35] and DAEM [36].

Estimation of the Beta vector π : After the inference of missing data associated with the users' taste matrix, we can first solve the posterior distributions over the Beta realizations $\{\pi_d\}_{d=1}^D$ used in the stick-breaking construction of Dirichlet Process. Under the principle in Eq. 19, we have

$$\ln q(\pi) \propto \ln \prod_{d=1}^D \pi_d^{\rho_d} (1 - \pi_d)^{\alpha + \sum_{j=d+1}^D \rho_j - 1}, \quad (25)$$

where $\rho_d = \sum_{i=1}^U q_i(d)$ represents the expected number of users coming from the d -th community. It is obvious that each element in the Beta vector is independent with each other and beta distributed. For π_d , the parameters of its posterior distribution is as follows:

$$\eta_{1,d} = \rho_d + 1, \quad \eta_{2,d} = \alpha + \sum_{j=d+1}^D \rho_j. \quad (26)$$

Thus, the expected value of the d -th stick length is $\frac{\rho_d + 1}{\sum_{j=d}^D \rho_j + 1 + \alpha}$, where α provides a basis term.

Estimation of the d -th community mean vector μ_d : According to the principle in Eq. 19, we can get that the posterior distribution over μ_d satisfies the following formula,

$$\begin{aligned} \ln q(\mu_d) \propto & -\frac{1}{2} \mu_d^T \left(\sum_{i=1}^U q_i(d) \langle \Sigma_d^{-1} \rangle + \frac{\langle \Sigma_d^{-1} \rangle}{\lambda_0} \right) \mu_d \\ & + \left(\sum_{i=1}^U q_i(d) \langle a_i^T \rangle \langle \Sigma_d^{-1} \rangle + \langle \mu_0^T \rangle \frac{\langle \Sigma_d^{-1} \rangle}{\lambda_0} \right) \mu_d. \end{aligned} \quad (27)$$

Similar to the inference of users' taste matrix A , the posterior distribution of μ_d is also Gaussian distributed with:

$$\begin{aligned} \langle \mu_d \rangle &= \frac{\lambda_0}{\lambda_0 + \rho_d} \sum_{i=1}^U q_i(d) \langle a_i \rangle + \frac{1}{\lambda_0 + \rho_d} \mu_0, \\ \langle \mu_d \mu_d^T \rangle &= \frac{\lambda_0}{\lambda_0 + \rho_d} \langle \Sigma_d^{-1} \rangle^{-1} + \langle \mu_d \rangle \langle \mu_d^T \rangle. \end{aligned} \quad (28)$$

It can be observed that the expectation of the d -th mean vector, $\langle \mu_d \rangle$, consists of two terms corresponding to the weighted mean of all users' taste vectors, and the prior knowledge to avoid the over-fitting problem respectively. Eq. 28 states that the properties of each community is fully determined by the users belonging to it.

Estimation of the d -th community covariance matrix Σ_d : Finally, we introduce our approach to inference the second-order statistics of the d community, denoted as the covariance matrix Σ_d . Following Eq. 19, we have,

$$\begin{aligned} \ln q(\Sigma_d | \mu_d) \propto & -\frac{1}{2} \text{tr} \sum_{i=1}^U q_i(d) (a_i - \mu_d)(a_i - \mu_d)^T \Sigma_d^{-1} \\ & - \frac{1}{2} \sum_{i=1}^U q_i(d) \ln |\Sigma_d| - \frac{1}{2} \text{tr}(W_0 \Sigma_d^{-1}) - \frac{\iota_0 + r + 1}{2} \ln |\Sigma_d| \\ & - \frac{1}{2} \left[\text{tr} \frac{(\mu_d - \mu_0)(\mu_d - \mu_0)^T}{\lambda_0} \Sigma_d^{-1} + \ln |\Sigma_d| \right]. \end{aligned} \quad (29)$$

Thus, the posterior distribution of Σ_d is also an inverse-Wishart distribution $i\mathcal{W}(W_d, \iota_d)$ with following parameters:

$$\begin{aligned} W_d &= \sum_{i=1}^U q_i(d) (\langle a_i a_i^T \rangle - 2\langle \mu_d \rangle \langle a_i^T \rangle + \langle \mu_d \mu_d^T \rangle) \\ &\quad + \frac{1}{\lambda_0} (\langle \mu_d \mu_d^T \rangle - 2\mu_0 \langle \mu_d^T \rangle + \mu_0 \mu_0^T) + W_0, \\ \iota_d &= \iota_0 + \rho_d, \end{aligned} \tag{30}$$

Thus, we can subsequently get

$$\begin{aligned} \langle \Sigma_d^{-1} \rangle &= \iota_d W_d^{-1}, \\ \langle \ln |\Sigma_d| \rangle &= \ln |\phi_d| - \psi\left(\frac{\iota_d}{2}\right) - r \ln(2). \end{aligned} \tag{31}$$

Algorithm 1: Batch Variational Bayesian (bVB) Inference for HeMF

Input: A rating matrix R observed on positions Ω , and the hyper-parameters Ξ

Output: The completed matrix \hat{R}

Initialization :

- 1: **while** not converge **do**
 - 2: Estimate the users' taste matrix A by Eq. 21
 - 3: Estimate the membership vector of users' taste Z
by Eq. 24
 - 4: Estimate the Beta vector π by Eq. 26
 - 5: Estimate the d -th community mean vector μ_d by
Eq. 28
 - 6: Estimate the d -th community covariance matrix Σ_d by Eq. 31
 - 7: **end while**
 - 8: **return** $\hat{R} = \langle A \rangle^T \langle B \rangle$
-

By iteratively updating formulas from Eq. 21 to 31, we develop an inference method for the proposed HeMF model, which can achieve at least a local optimal solution. The procedure is summarized in the Alg. 1. Noted that our inference method will degenerate to the famous Bayesian Probabilistic Matrix Factorization (BPMF) model [12] by setting $K = D = 1$. At each step in Alg. 1, we update the selected parameter with other parameters fixed. Thus, the lower bound of the marginal distribution is maximized successively in our method, whose convergence is naturally guaranteed.

5 Online Variational Inference

The method introduced in the last section requires to load the whole dataset for inference. However, when dealing with massive data, which becomes common nowadays, the memory space may be not enough for the bVB inference. To tackle this problem, we now develop an online variational inference method. We assume that a local optimal configuration has been achieved to approximate the marginal distribution of observing the rating matrix R_Ω at time t , and there are totally D user communities and K item groups. If new scores are observed at entire Ω' at time $t + 1$, the marginal distribution of observing the rating matrix $R_{\Omega \cup \Omega'}$ can be reformulated as:

$$\begin{aligned} p(R_{\Omega \cup \Omega'} | \Xi) &= \int dY d\Theta_{D+1, K+1} p(Y, \Theta_{D+1, K+1} | \Xi) \\ &\quad \cdot p(R_\Omega | Y, \Theta_{D+1, K+1}, \Xi) p(R_{\Omega'} | Y, \Theta_{D+1, K+1}, \Xi). \end{aligned} \tag{32}$$

It can be seen that the computation of marginal distribution for the newly observed rating matrix $R_{\Omega \cup \Omega'}$ consists of three parts, the prior distribution $p(Y, \Theta_{D+1, K+1} | \Xi)$, the likelihood of observing the rating matrix

R_Ω and the likelihood of observing the rating matrix $R_{\Omega'}$. In Eq. 32, we set $\Theta_{D+1,K+1}$ as the integral term, which provides our model the capability to generate new components on the fly. To approximate the marginal distribution $p(R_{\Omega \cup \Omega'} | \Xi)$, we assume that a new trial distribution $q^{t+1}(Y, \Theta_{D+1,K+1})$ can be found over the restricted function, as VB,

$$q^{t+1}(Y, \Theta_{D+1,K+1}) = q^{t+1}(Y) q^{t+1}(\pi) q^{t+1}(\omega) \cdot \prod_{d=1}^{D+1} q^{t+1}(\phi_d) \prod_{k=1}^{K+1} q^{t+1}(\psi_k). \quad (33)$$

Let θ_j denote the j -th parameter belonging to the set of parameters $Y \cup \Theta_{D+1,K+1}$. Given the new observations sampled on the positions Ω' , the variational approximation problem with respect to $q(\theta_j)$ is equivalent to

$$\begin{aligned} \ln p(R_{\Omega \cup \Omega'} | \Xi) &\geq \int q^{t+1}(\theta_j) \mathbb{E}_{\neq j} \ln p(R_\Omega | Y, \Theta_{D+1,K+1}, \Xi) d\theta_j \\ &+ \int q^{t+1}(\theta_j) \mathbb{E}_{\neq j} \ln p(R_{\Omega'} | Y, \Theta_{D+1,K+1}, \Xi) d\theta_j \\ &+ \int q^{t+1}(\theta_j) \mathbb{E}_{\neq j} \ln p(Y, \Theta_{D+1,K+1} | \Xi) d\theta_j \\ &- \int q^{t+1}(\theta_j) \ln q^{t+1}(\theta_j) d\theta_j. \end{aligned} \quad (34)$$

In this way, the close-form solution of $q(\theta_j)$ satisfies the following condition:

$$\begin{aligned} q^{t+1}(\theta_j) &\propto \exp[\mathbb{E}_{\neq j} \ln p(R_{\Omega'} | Y, \Theta_{D+1,K+1}, \Xi)] \\ &\quad \cdot \exp[\mathbb{E}_{\neq j} \ln p(R_\Omega, Y, \Theta_{D+1,K+1} | \Xi)] \\ &\propto q^t(\theta_j) \exp[\mathbb{E}_{\neq j} \ln p(R_{\Omega'} | Y, \Theta_{D+1,K+1}, \Xi)] \end{aligned} \quad (35)$$

Compared with the traditional Variational Bayesian updating formula (18), there is one additional term in our online inference algorithm, which introduces the influence of new observations into the updating formula. There are some related works proposed for deploying the principle of variational Bayesian for large-scale datasets, including stochastic variational inference (SVI¹) [25], stream variational inference (SVI²) [26], Population Variational Bayesian (pVB) [27] and sequential variation approximation (SVA) [28]. Compared with these seminal works, our method is different by using the Bayesian decomposition of the joint distribution, which is simple but efficient. Moreover, to our best knowledge, no previous online variational method has been applied to the DDP task.

As illustrated in Fig. 1, new entries can be observed in the rating matrix R as time goes on. Thus, the model should be updated sequentially. In this section, our focus mainly concentrates on the mathematical derivation of the parameters associated with the users' taste matrix A , as Sec. 4. Assuming we have achieved a local optimal solution for the observed rating matrix R_Ω , we should update the parameters by following formulas in each single pass with the newly observed entries Ω' .

Online estimation of the users' taste matrix A : As proved in Sec. 4, the posterior distribution over a_i is also a Gaussian distribution. Thus, under the principle of online variational Bayesian (oVB), as shown in Eq. 35, we derive the updating formulas over a_i for different cases as follows,

1 If the i -th user rates new items, we get

$$\begin{aligned} \langle a_i \rangle^{t+1} &= \langle a_i \rangle^t + \frac{1}{\sigma^2} \Sigma_{a_i}^t \Delta_{1,i}^{1+t} \\ &\quad - \Sigma_{a_i}^t (\sigma^2 \mathbf{I} + \Delta_{2,i}^{t+1} \Sigma_{a_i}^t)^{-1} \Delta_{2,i}^{t+1} (\langle a_i \rangle^t + \frac{1}{\sigma^2} \Sigma_{a_i}^t \Delta_{1,i}^{1+t}), \\ \Sigma_{a_i}^{t+1} &= \Sigma_{a_i}^t - \Sigma_{a_i}^t (\sigma^2 \mathbf{I} + \Delta_{2,i}^{t+1} \Sigma_{a_i}^t)^{-1} \Delta_{2,i}^{t+1} \Sigma_{a_i}^t. \end{aligned} \quad (36)$$

where $\Delta_{1,i}^{1+t} = \sum_{(i,j) \in \Omega'} r_{i,j}^{1+t} \langle b_j \rangle^t$, and $\Delta_{2,i}^{t+1} = \sum_{(i,j) \in \Omega'} \langle b_j b_j^T \rangle^t$ summarize the information of newly rated items. Σ_{a_i} denotes the covariance matrix of the posterior distribution over a_i , and it is equivalent to $\Sigma_{a_i} = \langle a_i a_i^T \rangle - \langle a_i \rangle \langle a_i \rangle^T$. As we can observe, the expected user's taste vector will be updated with the newly rated items.

- 2 If the i -th user changes his previous rating scores, we should update the user's taste by Eq. 36, but with $\Delta_{1,i}^{1+t} = \sum_{(i,j) \in \Omega'} (r_{i,j}^{t+1} - r_{i,j}^t) \langle b_j \rangle^t$ and $\Delta_{2,i}^{t+1}$ remaining the same form. It can be understood practically that the re-rated item will enhance its weight in describing the user's taste if this user shows higher preference to this object.
- 3 If someone registers a new account in the recommendation system and rates some items, the user's taste vector can be initialized by Eq. 21 directly, which solves the problem of *cold start* naturally.

Such updating problem has been previously tackled by Stochastic Gradient Descent (SGD) methods for matrix factorization. Starting with some initial value a_i^t , SGD refines the parameter value by iterating the stochastic difference equation as follows:

$$a_i^{t+1} = a^t - \epsilon^t \mathcal{L}'(a^t), \quad (37)$$

where $\epsilon^{[t]}$ denotes a sequence of decreasing step sizes and $-\mathcal{L}'(a^t)$ is the direction of steepest descent given the new observations. Compared with SGD, our method gets a local optimal solutions for $\langle a_i \rangle$ and $\Sigma_{a_i}^{t+1}$ at each iterative step. As shown empirically, it usually yields a faster convergence speed without the need of tuning the learning rate. Moreover, under the framework of Bayesian theory, we can integrate the prior knowledges into our model and avoid the overfitting problem by considering a family of models.

Online estimation of the membership vector Z : With Eq. 36, we can update the user's taste with new observations. Before updating the membership of user i , let

$$\begin{aligned} \langle \Sigma_0^{-1} \rangle &= \iota_0 W_0^{-1}, \\ \langle \ln |\Sigma_0| \rangle &= \ln |W_0| - \psi\left(\frac{\iota_0}{2}\right) - r \ln(2), \end{aligned} \quad (38)$$

which denote some initial parameters of the users' community. Thus, the probability of introducing a new component for $\langle a_i \rangle^{t+1}$ satisfies

$$\ln q(z_i = D + 1) \propto \gamma_i(D + 1) = \xi_{1,i}(0) + \xi_{2,i}(0), \quad (39)$$

where $\xi_{1,i}(0), \xi_{2,i}(0)$ share the same forms with Eq. 23. If we allow the online scheme to introduce a new component for every updating pass, there will be infinite components asymptotically, which is unnecessary. Empirically, $\gamma_i(D + 1)$ is negligible enough for most passes, which indicates that the updated user's tastes can be adequately explained by existing communities, and there is no need of new community. In practice, we set a small value ϵ and increase D to $D + 1$ only when $\gamma_i(D + 1) > \epsilon$. This simple strategy is very efficient in controlling the model size in practical applications.

Thus, the indicator variable can be updated as follows:

$$q_i(d) = \frac{\exp(\gamma_i(d))}{\sum_{t=1}^D \exp(\gamma_i(t))}, \quad (40)$$

which shares the same formulation with Eq. 22, besides the fact that one more community may be generated into the learning process on the fly.

To simplify our derivations, let

$$\begin{aligned} \vartheta_1^{t+1}(d) &= \sum_{i \in \Omega'} q_i^{t+1}(d) - q_i^t(d), \\ \vartheta_2^{t+1}(d) &= \sum_{i \in \Omega'} q_i^{t+1}(d) \langle a_i \rangle^{t+1} - q_i^t(d) \langle a_i \rangle^t, \\ \vartheta_3^{t+1}(d) &= \sum_{i \in \Omega'} q_i^{t+1}(d) \langle a_i a_i^T \rangle^{t+1} - q_i^t(d) \langle a_i a_i^T \rangle^t, \end{aligned} \quad (41)$$

which denotes the updates of the missing data for the i -th user's taste. For a newly registered user, we have $q_i(d) = 0, \forall i \in \{1, \dots, D\}$. Next, we will move on to introduce the updating of model parameters, Θ .

Online estimation of the Beta vector π : It has been known that the posterior distribution over the d -th entry in the vector π is also a beta distribution with form $\text{Beta}(\eta_{1,d}, \eta_{2,d})$. Under the principle of oVB, the beta distribution parameters can be updated as follows,

$$\begin{aligned}\eta_{1,d}^{t+1} &= \eta_{1,d}^t + \vartheta_1^{t+1}(d), \\ \eta_{2,d}^{t+1} &= \eta_{2,d}^t + \sum_{j=d+1}^k \vartheta_1^{t+1}(j).\end{aligned}\tag{42}$$

It can be found that the posterior distribution parameters are adjusted by the updating term of membership vector.

Online estimation of the d -th community mean vector μ_d : Under the principle in Eq. 35, the posterior of μ_d is also a Gaussian distribution with parameters:

$$\langle \mu_d \rangle^{t+1} = \langle \mu_d \rangle^t - \frac{\vartheta_1^{t+1}(d) \langle \mu_d \rangle^t - \lambda_0 \vartheta_2^{t+1}(d)}{\lambda_0 + \rho_d^{t+1}},\tag{43}$$

with $\rho_d^{t+1} = \rho_d^t + \vartheta_1^{t+1}(d)$, denoting the updated expected number of users in the d -th community. Additionally, we have

$$\langle \mu_d \mu_d^T \rangle^{t+1} = \frac{\lambda_0}{\lambda_0 + \rho_d^{t+1}} \langle \Sigma_d^{-1} \rangle^{t-1} + \langle \mu_d \rangle^{t+1} \langle \mu_d^T \rangle^{t+1}.\tag{44}$$

To simplify the formulations for sequential estimation of Σ_d , let

$$\begin{aligned}\vartheta_4^{t+1}(d) &= \langle \mu_d \rangle^{t+1} - \langle \mu_d \rangle^t, \\ \vartheta_5^{t+1}(d) &= \langle \mu_d \mu_d^T \rangle^{t+1} - \langle \mu_d \mu_d^T \rangle^t, \\ \vartheta_6^{t+1}(d) &= \langle \mu_d \rangle^{t+1} \sum_{i \in \Omega'} \langle a_i^T \rangle^{t+1} - \langle \mu_d \rangle^t \sum_{i \in \Omega'} \langle a_i^T \rangle^t,\end{aligned}\tag{45}$$

which represent the updates of the d -th community mean vector.

Online estimation of the d -th community covariance matrix Σ_d : The posterior distribution of Σ_d is also an inverse-Wishart distribution, $i\mathcal{W}(W_d^{t+1}, \iota_d^{t+1})$ with:

$$\begin{aligned}W_d^{t+1} &= W_d^t - \frac{2}{\lambda_0} \mu_0 \vartheta_4^{t+1}(d) - 2\vartheta_6^{t+1}(d) \\ &\quad + (U + \frac{1}{\lambda_0}) \vartheta_5^{t+1}(d) + \vartheta_3^{t+1}(d), \\ \iota_d^{t+1} &= \iota_d^t + \vartheta_1^{t+1}(d),\end{aligned}\tag{46}$$

Thus, we can subsequently get

$$\begin{aligned}\langle \Sigma_d^{-1} \rangle^{t+1} &= W_d^{t+1} \iota_d^{t+1}, \\ \langle \ln |\Sigma_d| \rangle^{t+1} &= \ln |W_d^{t+1}| - \psi\left(\frac{\iota_d^{t+1}}{2}\right) - r \ln(2).\end{aligned}\tag{47}$$

In summary, given the newly observed entries, we can update the proposed HeMF model by successively executing the formulas from Eq. 36 to Eq. 47. The detailed procedure can be found in Alg. 2. It can be observed that the whole sequential updating procedures are triggered by the update of user's taste. Under the principle of Variational Bayesian, we update the selected user's taste with other parameters fixed.

Algorithm 2: online Variational Bayesian (oVB) inference for HeMF

Input: The new entries observed on positions Ω' , the hyper-parameters Ξ , and the local optimal solution for the posterior distribution $q^t(Y, \Theta)$

Output: The completed matrix \hat{R}

Initialization :

- 1: Update the users' taste matrix A by Eq. 36
 - 2: Update the membership vector of users' taste Z by Eq. 40
 - 3: Update the Beta vector π by Eq. 42
 - 4: Update the d -th community mean vector μ_d by Eq. 43
 - 5: Update the d -th community covariance matrix Σ_d by Eq. 47
 - 6: **return** $\hat{R} = \langle A \rangle^T \langle B \rangle$
-




Figure 3: The evolution of users' graph, where two users will be linked if they rate the same items. (a), (b) and (c) denote the users' graphs at time t , $t + 1$ and $t + 2$ respectively.

5.1 Notes about oVB

The previous section discussed the procedure to update the given model configuration when there are newly observed entries. Actually, the online updating procedure, Alg. 2, should be triggered every time when some parameters change. To help understand the proposed oVB, this section comprehensively explains the updating behaviors of parameters when there are the newly observed rating entries.

As illustrated in Fig. 3, we have constructed a graph for the sparsely observed matrix shown in Fig. 1, where two users are linked when they have rated same items. It can be observed that at time t , only the 3rd user and the 8th user are linked. At time $t + 1$, the 6th user rated the 5th item, thus a link between the 4th user and the 6th user is established. At time $t + 2$, it can be found that there are no more isolated nodes in the graph, resulting to a connected graph. With these graphs as an example, we will explain how our online learning algorithm behaves.

At time $t + 1$, the 6th user rated the 5th item, the 6th user's taste vector a_6 should be updated with Eq. 36, which states that the expected taste vector would be updated by a term with respect to the newly rated item b_5 . At the same time, the property vector of the 5th item should also be changed due to the 6th user. It can be found that the taste vector of 4th user is also determined by his/her rated items, including I_5 and I_7 , and the updating of the 6th user's taste vector would cause the updating of the 4th user via their commonly rated item, I_5 . At time $t + 2$, it can be found that each node in the graph requires to be re-estimated when there is any newly observed rating. And the model parameters should be updated based on the newly estimated taste and property vectors accordingly.

Practically, we can divide a massive dataset into small ones, and import them into Alg. 2 sequentially. As stated previously, our method can provide a better solution to explain the entire massive dataset after each updating pass.

6 Empirical Variational Method

For both bVB and oVB, a significant issue is that the hyper-parameters may be unknown empirically. This section addresses the issue and develops the empirical Variational Method [37] for HeMF, which updates the hyper-parameters for each iteration. This method can further optimize the lower bound of $\mathcal{L}_{q(Y, \Theta)}(R_\Omega | \Xi)$ resulting to better performance. By taking the derivative of the lower bound with respect to each hyper-parameter, we derive the updating scheme as follows.

Optimizing the noise variance σ^2 : By setting the derivative $\frac{\partial \mathcal{L}(\sigma^2)}{\partial \sigma^2} = 0$, we get the optimal solution for σ^2 as follows,

$$\sigma^2 = \frac{1}{|\Omega|} \sum_{(i,j) \in \Omega} r_{i,j}^T r_{i,j} - 2 \langle b_j^T \rangle \langle a_i \rangle r_{i,j} + \text{Tr} \langle b_j b_j^T \rangle \langle a_i a_i^T \rangle, \quad (48)$$

where $|\Omega|$ denotes the number of observed entries. In the online setting, the noise variance σ^2 can be updated with

$$\begin{aligned} \sigma^{2^{t+1}} &= \frac{|\Omega|}{|\Omega \cup \Omega'|} \sigma^{2^t} \\ &+ \frac{1}{|\Omega \cup \Omega'|} \sum_{(i,j) \in \Omega'} r_{i,j}^T r_{i,j} - 2 \langle b_j^T \rangle \langle a_i \rangle r_{i,j} + \text{Tr} \langle b_j b_j^T \rangle \langle a_i a_i^T \rangle. \end{aligned} \quad (49)$$

Optimizing the parameter over Beta distributions α and β : The derivative of $\mathcal{L}(\alpha)$ with respect to α can be written as follows

$$\begin{aligned} \mathcal{L}'(\alpha) &= \sum_{d=1}^D [\psi(\eta_{2,d}) - \psi(\eta_{1,d} + \eta_{2,d})] \\ &+ (\alpha - 1)D[\ln \Gamma(\alpha) + \psi(\alpha)] \end{aligned} \quad (50)$$

The new α can be obtained by the gradient method as:

$$\alpha^{new} = \alpha^{old} - \epsilon_\alpha \mathcal{L}'(\alpha), \quad (51)$$

where ϵ_α is the learning rate used for updating α . By altering the notions, we can get the updating formula for β in the same way.

Optimizing the parameter over μ_0 : The lower bound function with respect to μ_0 can be reformulated as follows:

$$\mathcal{L}(\mu_0) \propto \mu_0^T \langle \Sigma^{-1} \rangle \langle \mu \rangle - \frac{1}{2} \mu_0^T \langle \Sigma^{-1} \rangle \mu_0. \quad (52)$$

And the optimal solution for μ_0 is

$$\mu_0 = \langle \mu \rangle = \frac{1}{D} \sum_{d=1}^D \mu_d. \quad (53)$$

The optimal solution for ν_0 can be achieve alternatively as $\nu_0 = \langle \nu \rangle$.

Optimization the parameter over λ_0 : λ_0 is a hyperparameter controlling the diversity of the prior distribution of μ_0 . By setting $\frac{\partial \mathcal{L}(\lambda_0)}{\partial \lambda_0} = 0$, we can solve the optimal solution for λ_0 with,

$$\begin{aligned} \lambda_0 &= \frac{\sum_{d=1}^D (\text{tr} \langle \mu_d \mu_d^T \rangle \langle \Sigma_d^{-1} \rangle - 2 \mu_0^T \langle \Sigma_d^{-1} \rangle \langle \mu_d \rangle + \mu_0^T \langle \Sigma_d^{-1} \rangle \mu_0)}{L(D + K)} \\ &+ \frac{\sum_{k=1}^K (\text{tr} \langle \nu_k \nu_k^T \rangle \langle \Pi_k^{-1} \rangle - 2 \nu_0^T \langle \Pi_k^{-1} \rangle \langle \nu_k \rangle + \nu_0^T \langle \Pi_k^{-1} \rangle \nu_0)}{L(D + K)}, \end{aligned} \quad (54)$$

which summarizes the diversity of the model parameters $\{\mu_d\}_{d=1}^D$ and $\{\nu_k\}_{k=1}^K$.

Optimizing the parameter over W_0 and ι_0 : Finally, we derive the updating formulas for hyper-parameters controlling the inverse-wishart distribution. As shown in [38], the detailed formulation of inverse-wishart distribution can be written as,

$$\begin{aligned} \ln i\mathcal{W}(\Sigma_d|W_0, \iota_0) &= -0.5\text{tr}(W_0\Sigma_d^{-1}) - \frac{\iota_0 + L + 1}{2} \ln |\Sigma_d| \\ &+ \frac{\iota_0}{2} \ln |W_0| - \frac{\iota_0 L}{2} \ln 2 - \ln \Gamma_L(\frac{\iota_0}{2}) \end{aligned} \quad (55)$$

Thus, by taking the derivative of \mathcal{L} with respect to W_0 and ι_0 , we can get the updating formulas as follows,

$$\begin{aligned} W_0 &= (D + K)\iota_0 \left(\sum_{d=1}^D \Sigma_d^{-1} + \sum_{k=1}^K \Pi_k^{-1} \right)^{-1} \\ \iota_0^{new} &= \iota_0^{old} - \epsilon_{\iota_0} \mathcal{L}'(\iota_0), \end{aligned} \quad (56)$$

with

$$\begin{aligned} \mathcal{L}'(\iota_0) &= \frac{(D + K)}{2} [\ln |W_0| - L \ln 2 - \psi_L(\frac{\iota_0}{2})] \\ &- \frac{1}{2} \left(\sum_{d=1}^D \ln |\Sigma_d| + \sum_{k=1}^K \ln |\Pi_k| \right) \end{aligned} \quad (57)$$

The detailed procedures of optimizing the hyper-parameters can be found in Alg. 3. Since all of the hyper-parameters are updated by the gradient method, our empirical variational methods are theoretically guaranteed to converge.

Algorithm 3: empirical Variational Bayesian (eVB) inference for HeMF

Input: The rating R observed on positions Ω

Output: The completed matrix \hat{R}

Initialisation :

- 1: **while** non converge **do**
 - 2: The estimating procedures in Alg. 1 or Alg. 2.
 - 3: Optimize the noise variance σ^2 by Eq. 48 or 49.
 - 4: Optimize the hyperparameter of Beta distributions α and β by Eq. 51
 - 5: Optimize the hyperparameter μ_0, ν_0 by Eq. 53
 - 6: Optimize the hyperparameter λ_0 by Eq. 54
 - 7: Optimize the hyperparameter W_0 and ι_0 by Eq. 56
 - 8: **end while**
 - 9: **return** $\hat{R} = \langle A \rangle^T \langle B \rangle$
-

7 Experiments and Discussions

To evaluate the performance of the proposed methods, we apply them on 4 benchmark movie rating collaborative filtering datasets, i.e., the Netflix Prize dataset¹, the EachMovie dataset², the 1M and Latest MovieLens Datasets³. We take two state-of-the-art methods as the baseline approaches, SGD and M³F [17], whose source codes can be downloaded from the website⁴. For a fair comparison, all experiments are

¹<http://www.netflixprize.com/>

²<http://grouplens.org/datasets/eachmovie/>

³<http://grouplens.org/datasets/movielens/>

⁴<https://code.google.com/p/m3f/>

1M MovieLens		
Method	Training error	Testing error
M ³ F-TIB(1,1,30)	0.8029	0.8491
M ³ F-TIB(2,1,40)	0.7978	0.8478
M ³ F-TIF(1,2,30)	0.8015	0.8484
M ³ F-TIF(2,2,40)	0.8012	0.8489
HeMF-bVB	0.7723	0.8311
EachMovie		
Method	Training error	Testing error
M ³ F-TIB(1,2,40)	0.5924	1.0993
M ³ F-TIB(2,1,40)	0.6145	1.0920
M ³ F-TIF(1,1,30)	0.7863	1.1021
M ³ F-TIF(1,2,30)	0.7846	1.1040
HeMF-bVB	0.8835	1.0845

Table 1: 1M MovieLens and EachMovie RMSE score for two M³F models with optimal configurations, All scores are averaged across 10 standardized cross-validation splits. Parentheses indicate topic counts as well as the initial hidden dimensionality for M³F models (K^U, K^M, D). Best results for each dataset have been boldened.

conducted on an Intel Core i7 920 2.67GHz CPU with 12 G RAM. For SGD and M³F, we use the default parameters reported in their papers, and set the same initial configurations with our methods.

A standard cross-validation technique [39] is applied to estimate the performance of each approach. 90% entries are randomly selected as the training set, while the left 10% entries are treated as the unobserved dyadic for testing. To compare the online learning performance of oVB with SGD, we further design a series of online learning experiments. After each standard cross-validation split, we further divide the training set into several smaller ones, which will then be sub-sequentially and repeatedly imported into the online learning procedures, oVB and SGD. As most researchers do, we evaluate the performance of all methods by RMSE of testing dataset, which is defined as:

$$RMSE = \sqrt{\frac{1}{\sum_{i,j \in \Omega} 1} \sum_{i,j \in \Omega} (r_{ij} - \hat{r}_{ij})^2}, \quad (58)$$

where \hat{r}_{ij} denotes the predicted rating value. It is obvious that a smaller RMSE value indicates better performance.

In all the experiments conducted in this paper, we apply an empirical variational Bayesian (eVB) updating procedures after both batch variational Bayesian (bVB) and online variational Bayesian (oVB). During the implementation of oVB, we observe that some components introduced at early sequential pass would become useless. We thus introduce a mechanism to merge similar communities. The only required parameters for our methods are the learning rates used in the eVB procedures, which is set to be 0.001 simply in all the experiments.

7.1 1M MovieLens and EachMovie Datasets

We first evaluate our models on the smaller datasets, 1M MovieLens and EachMovie, which can be effectively processed by bVB and M³F models. The EachMovie dataset contains 2.8 million ratings in $\{1, \dots, 6\}$ distributed across 1648 movies and 74424 users. The 1M movieLens dataset has 6040 users, 3952 movies, and 1 million ratings in $\{1, \dots, 5\}$. Following the "weak generalization" ratings prediction experiment in [13, 17], for each user in the training set, we withhold a single rating for the test set. All reported results are averaged over the same 10 random train-test cross-validate splits used in [13, 17].

Table 1 reports the predictive performance of the proposed Heterogeneous Matrix Factorization (HeMF) Model inferred by bVB method. To achieve the best performance of competitive methods, we implement M³F [17] with a variety of factor dimensionality and topic counts. The 4 optimal configurations of M³F has been selected for our comparisons. All results are reported as the normalized mean average error.




Figure 4: The testing RMSE of M³F and HeMF on 1M MovieLens dataset.

It can be observed that our proposed HeMF model inferred by bVB systematically outperforms the current state-of-the-art models, M³F models, on both the 1M MovieLens and EachMovie datasets. Note that, for EachMovie datasets, it can be found that all M³F methods achieve better training RMSE results, but with worse testing performance, which indicates that M³F suffers from the overfitting problem while the proposed HeMF model well avoid it. As shown in Fig. 4, even starting with a worse initial configuration, the proposed HeMF model can converges to a better optimal solution using less iterative steps.

Since the initial hidden dimensionality of each component in HeMF is set as 20, the results indicate that the assumption of mixed membership offers greater predictive power than simply varying the dimensionality and topic counts in M³F models. Moreover, with the help of Dirichlet Process, the component of groups and the corresponding dimensionality can be determined automatically during the leaning process.

We further compare our online variational Bayesian (oVB) inference approach with the most recent sequential learning of Matrix Factorization, SGD. Starting with the same initial configuration, both SGD and oVB are fed with the same sequential observations, which contains 30 samples each pass in this experiment. Fig. 5 illustrates the predictive performance of our sequential variational approach and the stochastic gradient descent method. It can be observed that the solution given by SGD always vibrates around an unsatisfied point, while the testing RMSE given by oVB decreases continuously and finally converges to a local optimal solution.

To investigate the performance of the proposed oVB under different conditions, we further test it on sequential datasets with different size. As confirmed by Fig. 6, oVB converges much faster and gets better performance when fed with more samples in each sequential pass.




Figure 5: The testing RMSE of oVB and SGD on 1M MovieLens dataset.

7.2 Latest MovieLens Dataset

We further test our methods on a larger dataset, the Latest MovieLens Dataset, which contains only 21,622,187 ratings distributed across 234,934 users and 30,106 movies. Thus, only 0.3% entries are observed in this matrix, which is much more challenging for the DDP task. Since M³F-TIB model with the configuration (2, 1, 40) achieves the best performance among the M³F family, we set it as our baseline.

The performance of each iterative step conducted in M³F-TIB and HeMF has been illustrated in Fig. 7. It can be observed that our method has much better performance with respect to RMSE value when starting with the same initial configuration. With the size of each sequential dataset as 60, the performance of oVB can be found in Fig. 8, which shows that oVB can achieve better results when given much richer data. When fed with more sequential data, the performance of traditional SGD becomes worse, which indicates that the SGD suffers from serious overfitting problem, which is the same as results on smaller datasets.

7.3 Netflix Dataset

We further implement all methods on the famous Netflix Prize dataset, which contains 100 million ratings in $\{1, \dots, 5\}$ distributed across 17,770 movies and 480,189 users. Because the public evaluation on the unobserved ratings for dyad Netflix Prize dataset is no longer available, we apply a standard cross-validation split on the observed ratings of Netflix Prize dataset, with the same setting on 1M MovieLens and EachMovie Datasets. This section sets M³F-TIB model with the configuration (2, 1, 200) as our baseline.

Fig. 9 and 10 show the experimental results. It can be observed that HeMF can achieve better results, while the performance of TIB model get even worse with more iterative steps, which shows that the HeMF has better predictive power.

Fig. 9 illustrates the online performance of oVB and SGD. It can be observed that the proposed SVA achieves better testing RMSE, while SGD gives an unstable estimation, and even get worse when fed with




Figure 6: The testing RMSE of oVB, on 1M MovieLens dataset, when fed with samples of different size in each sequential pass.

more samples.

8 Conclusions

In this work, we developed a novel Bayesian dyadic data prediction model which integrates the complementary approaches of discrete mixed membership modeling and continuous latent factor modeling, which successfully accounts the heterogeneous property of users and their interaction in the practical recommendation systems. Two variational methods were derived to solve the proposed DDP task, one is the batch variational approximation, another is the online variational method for large-scale DDP. The performance of our methods were evaluated on real datasets including EachMovie, MovieLens and Netflix Prize. On each dataset, we found that our HeMF model achieved superior performance compared with the M³F models. Specially, our online learning method oVB shows significant improvement compared with SGD, not only on the estimation accuracy but also on its robustness. It can be indicated that the proposed HeMF model with the inference methods can be a better candidate to cope with real-world DDP tasks, especially the application of recommendation systems. And the proposed online learning algorithm can well cope with the large-scale DDP problem in real-world applications.

References

- [1] T. Hofmann, J. Puzicha, and M. I. Jordan, “Learning from dyadic data,” *Advances in neural information processing systems*, pp. 466–472, 1999.




Figure 7: The testing RMSE of HeMF and TIB on Latest MovieLens dataset.

- [2] E. J. Candès and T. Tao, “The power of convex relaxation: Near-optimal matrix completion,” *Information Theory, IEEE Transactions on*, vol. 56, no. 5, pp. 2053–2080, 2010.
- [3] N. Natarajan and I. S. Dhillon, “Inductive matrix completion for predicting gene–disease associations,” *Bioinformatics*, vol. 30, no. 12, pp. i60–i68, 2014.
- [4] X. Sui, T.-H. Lee, J. J. Whang, B. Savas, S. Jain, K. Pingali, and I. Dhillon, “Parallel clustered low-rank approximation of graphs and its application to link prediction,” in *Languages and Compilers for Parallel Computing*. Springer, 2013, pp. 76–95.
- [5] A. Pascual-Montano, J. M. Carazo, K. Kochi, D. Lehmann, and R. D. Pascual-Marqui, “Nonsmooth nonnegative matrix factorization (nsnmf),” *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 28, no. 3, pp. 403–415, 2006.
- [6] R. Sandler and M. Lindenbaum, “Nonnegative matrix factorization with earth mover’s distance metric for image analysis,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 33, no. 8, pp. 1590–1602, 2011.
- [7] H. Liu, Z. Wu, X. Li, D. Cai, and T. S. Huang, “Constrained nonnegative matrix factorization for image representation,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 34, no. 7, pp. 1299–1311, 2012.
- [8] M. Zitnik and B. Zupan, “Data fusion by matrix factorization,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 37, no. 1, pp. 41–53, 2015.




Figure 8: The testing RMSE of oVB and SGD on Latest MovieLens dataset, when feed with 60 samples in each sequential pass.

- [9] N. Gillis, S. Vavasis *et al.*, “Fast and robust recursive algorithms for separable nonnegative matrix factorization,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 36, no. 4, pp. 698–714, 2014.
- [10] E. J. Candès and B. Recht, “Exact matrix completion via convex optimization,” *Foundations of Computational mathematics*, vol. 9, no. 6, pp. 717–772, 2009.
- [11] E. J. Candès and Y. Plan, “Matrix completion with noise,” *Proceedings of the IEEE*, vol. 98, no. 6, pp. 925–936, 2010.
- [12] T. Hofmann, “Probabilistic latent semantic indexing,” in *Proceedings of the 22nd annual international ACM SIGIR conference on Research and development in information retrieval*. ACM, 1999, pp. 50–57.
- [13] B. Marlin and R. S. Zemel, “The multiple multiplicative factor model for collaborative filtering,” in *Proceedings of the twenty-first international conference on Machine learning*. ACM, 2004, p. 73.
- [14] R. Salakhutdinov and A. Mnih, “Bayesian probabilistic matrix factorization using markov chain monte carlo,” in *Proceedings of the 25th international conference on Machine learning*. ACM, 2008, pp. 880–887.
- [15] E. M. Airoldi, D. M. Blei, S. E. Fienberg, and E. P. Xing, “Mixed membership stochastic blockmodels,” in *Advances in Neural Information Processing Systems*, 2009, pp. 33–40.
- [16] I. Porteous, E. Bart, and M. Welling, “Multi-hdp: A non parametric bayesian model for tensor factorization,” in *Aaai*, vol. 8, 2008, pp. 1487–1490.




Figure 9: The testing RMSE of HeMF and TIB on Netflix dataset.

- [17] L. W. Mackey, D. Weiss, and M. I. Jordan, “Mixed membership matrix factorization,” in *Proceedings of the 27th international conference on machine learning (ICML-10)*, 2010, pp. 711–718.
- [18] Y.-D. Kim and S. Choi, “Scalable variational bayesian matrix factorization with side information,” in *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, Reykjavik, Iceland, 2014.
- [19] S. Park, Y.-D. Kim, and S. Choi, “Hierarchical bayesian matrix factorization with side information,” in *Proceedings of the Twenty-Third international joint conference on Artificial Intelligence*. AAAI Press, 2013, pp. 1593–1599.
- [20] R. P. Adams, G. E. Dahl, and I. Murray, “Incorporating side information in probabilistic matrix factorization with gaussian processes,” *arXiv preprint arXiv:1003.4944*, 2010.
- [21] I. Porteous, A. U. Asuncion, and M. Welling, “Bayesian matrix factorization with side information and dirichlet process mixtures.” in *AAAI*, 2010.
- [22] H. Attias, “A variational bayesian framework for graphical models,” *Advances in neural information processing systems*, vol. 12, no. 1-2, pp. 209–215, 2000.
- [23] A. Mnih and R. Salakhutdinov, “Probabilistic matrix factorization,” in *Advances in neural information processing systems*, 2007, pp. 1257–1264.
- [24] B. K. L. T. T. S. C. Jaya Kawale, Hung Bui, “Efficient thompson sampling for online matrix-factorization recommendation,” in *Advances in Neural Information Processing Systems*, 2015.
- [25] T. Broderick, N. Boyd, A. Wibisono, A. C. Wilson, and M. I. Jordan, “Streaming variational bayes,” in *Advances in Neural Information Processing Systems*, 2013, pp. 1727–1735.




Figure 10: The testing RMSE of oVB and SGD on Netflix dataset, when feed with 60 samples in each sequential pass.

- [26] A. Tank, N. J. Foti, and E. B. Fox, “Streaming variational inference for bayesian nonparametric mixture models,” *arXiv preprint arXiv:1412.0694*, 2014.
- [27] J. McInerney, R. Ranganath, and D. Blei, “The population posterior and bayesian modeling on streams,” *arXiv preprint arXiv:1507.05253v2*, 2015.
- [28] D. Lin, “Online learning of nonparametric mixture models via sequential variational approximation,” in *Advances in Neural Information Processing Systems*, 2013, pp. 395–403.
- [29] T. S. Ferguson, “A bayesian analysis of some nonparametric problems,” *The annals of statistics*, pp. 209–230, 1973.
- [30] J. Sethuraman, “A constructive definition of dirichlet priors,” DTIC Document, Tech. Rep., 1991.
- [31] B. Recht, M. Fazel, and P. A. Parrilo, “Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization,” *SIAM review*, vol. 52, no. 3, pp. 471–501, 2010.
- [32] G. Parisi, *Statistical Field Theory*. Addison Wesley, 1988.
- [33] M. Abramowitz and I. A. Stegun, *Handbook of mathematical functions: with formulas, graphs, and mathematical tables*. Courier Corporation, 1964, no. 55.
- [34] A. P. Dempster, N. M. Laird, and D. B. Rubin, “Maximum likelihood from incomplete data via the em algorithm,” *Journal of the royal statistical society. Series B (methodological)*, pp. 1–38, 1977.
- [35] L. Xu, “Byy harmony learning, structural rpcl, and topological self-organizing on mixture models,” *Neural Networks*, vol. 15, no. 8, pp. 1125–1151, 2002.

- [36] N. Ueda and R. Nakano, “Deterministic annealing em algorithm,” *Neural Networks*, vol. 11, no. 2, pp. 271–282, 1998.
- [37] C. M. Bishop, *Pattern recognition and machine learning*. springer, 2006.
- [38] L. Haff, “An identity for the wishart distribution with applications,” *Journal of Multivariate Analysis*, vol. 9, no. 4, pp. 531–544, 1979.
- [39] Y. Bengio and Y. Grandvalet, “No unbiased estimator of the variance of k-fold cross-validation,” *The Journal of Machine Learning Research*, vol. 5, pp. 1089–1105, 2004.